

Volumetric three-dimensional display system with rasterization hardware

Gregg Favalora*, Rick K. Dorval, Deirdre M. Hall, Michael Giovinco, Joshua Napoli

Actuality Systems, Inc., 315 Main Street, Suite 102, Reading, MA 01867

ABSTRACT

An 8-color multiplanar volumetric display is being developed by Actuality Systems, Inc. It will be capable of utilizing an image volume greater than 90 million voxels, which we believe is the greatest utilizable voxel set of any volumetric display constructed to date. The display is designed to be used for molecular visualization, mechanical CAD, e-commerce, entertainment, and medical imaging. As such, it contains a new graphics processing architecture, novel high-performance line-drawing algorithms, and an API similar to a current standard.

Three-dimensional imagery is created by projecting a series of 2-D bitmaps ("image slices") onto a diffuse screen that rotates at 600 rpm. Persistence of vision fuses the slices into a volume-filling 3-D image. A modified three-panel Texas Instruments projector provides slices at approximately 4 kHz, resulting in 8-color 3-D imagery comprised of roughly 200 radially-disposed slices which are updated at 20 Hz. Each slice has a resolution of 768 by 768 pixels, subtending 10 inches. The display electronics includes a custom rasterization architecture which converts the user's 3-D geometry data into image slices, as well as 6 Gbits of DDR SDRAM graphics memory.

Keywords: 3-D display, volumetric, autostereoscopic, spatial light modulator, pharmaceutical, MCAD, entertainment

1. INTRODUCTION

Actuality Systems, Inc. was founded in 1997 with the goal of empowering knowledge workers to get as much use as possible out of their visual information. This paper describes the research done to date in the creation of a volumetric three-dimensional display system which is applicable to mechanical CAD¹, molecular visualization², medical imaging, entertainment, and e-commerce. The technology outlined here was designed with an eye towards commercialization from the outset.

1.1 Introduction to Volumetric Displays

Volumetric displays are a class of three-dimensional display technology which produces volume-filling imagery. As defined by Barry Blundell and Adam Schwarz, "[a] volumetric display device permits the generation, absorption, or scattering of visible radiation from a set of localized and specified regions within a physical volume."³ Most, if not all, volumetric displays are *autostereoscopic*; that is, they produce imagery that appears three-dimensional without the use of additional eyewear. The volumetric analogue to pixels are called *voxels*, short for volume elements (or volume pixels).

Many volumetric displays create three-dimensional imagery by employing spatio-temporal multiplexing in emitting or scattering light from a range of locations within that volume. That is, a smaller number of light-generating devices (lasers, projector pixels, etc.) are often run at a higher frequency than the overall volumetric refresh rate, and the light is typically imaged onto a surface undergoing periodic motion. For instance, in 1960, Richard Ketchpel filed a U.S. patent application⁴ for a volumetric display that uses a cathode ray gun to excite emissive regions on a rotating phosphored disk. Persistence of vision is used to integrate the emissions into a volume-filling whole. Vector-scanned 3-D displays such as Ketchpel's have been built in a variety of sizes and methodologies.^{5,6} However, because their beam steering and modulation components have limited bandwidth, vector-scanned displays have a very low percentage (historically on order of 1%) of *total usable* voxels per image despite a high *addressable* resolution.

In contrast, the display that will be described here is a *raster-scanned* or *bitmapped* 3-D display, since the user has the opportunity to use up to 100% of the addressable resolution in creating an image. The relationship between vector- and raster-scanned displays is similar to the relationship between oscilloscopes and high-resolution CRTs.

* Correspondence: Email: <favalora@actuality-systems.com>; WWW: <http://www.actuality-systems.com>; Telephone: (781) 944-8941.

It is beyond the scope of this paper to give a review of volumetric display technologies. The authors point the interested reader to a variety of summaries, such as Michael Halle's excellent survey⁷ of autostereoscopic displays, Blundell and Schwarz's book referenced above, and a recent review⁸ published by the FELIX team in Germany. They together reveal a rich 3-D display development history that also includes "solid-state" displays^{9,10,11} and varifocal mirror displays¹² not discussed here.

1.2 Related Prior Research

The research presented here is inspired by work done in the late 1950s and early 1960s, in which 3-D imagery is perceived via the visual integration of a series of images projected onto a rotating screen, wherein a fixed 2-D projector illuminates a diffuse rotating screen after being relayed by several mirrors that rotate with the screen. In this way, incoherent light sources can be used to form the imagery, since the path length from the projector to the screen is not a function of the screen's position. For instance, in 1958, Max Hirsch filed a patent application¹³ for his "generescope," in which imagery formed on the surface of a cathode ray tube is focused onto a periscope-like arrangement that images onto a rear-projection screen. An illustration from his U.S. patent is shown in Figure 1. Note that the CRT, mirrors, and screen rotate in unison.

A similar display which utilizes a *fixed* CRT and optical system but which projects onto a rotating mirror, prism, and screen system was developed by ITT Laboratories¹⁴ and described in an article in 1960. A consequence of illuminating a rotating assembly with a fixed screen is that the CRT's image rotates in the plane of the projection screen as the projection screen rotates. Robert Batchko describes a similar system that uses a vector-scanned laser illumination source.¹⁵

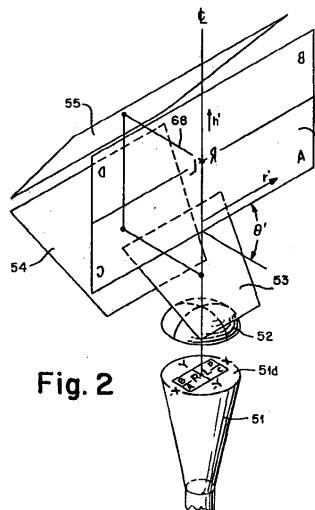


Fig. 2

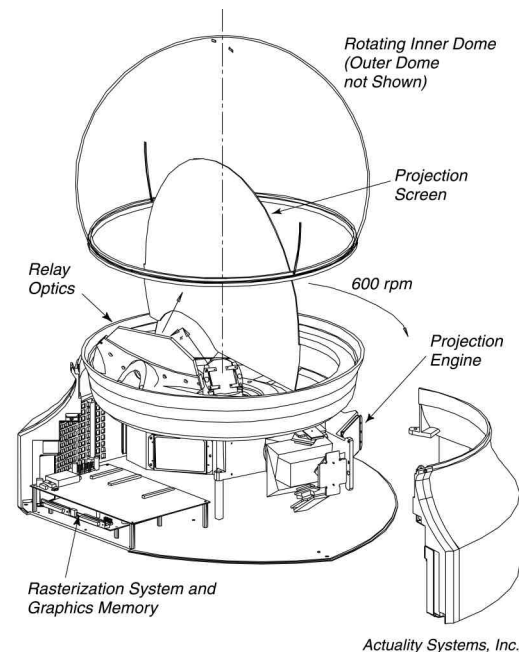


Figure 1. The "generescope." From M. Hirsch's U.S. Pat. 2,967,905

Figure 2. Actuality's 3-D display. Fixed outer dome, graphics memory, fan, and cabling omitted for simplicity.

Che-Chih Tsao *et al.* describe a volumetric display that uses a k-mirror system which rotates at half of the screen's rotational frequency to keep the projected image fixed on the surface of the screen.^{16,17}

Despite the improvements of volumetric display technology, none have achieved widespread commercialization. The authors attribute this to the historically limited resolution of these displays, as well as their previous difficulty in using them to augment existing visualization workflows. We therefore intend to produce a complete display system which includes a display of volumetric resolution at least 10x greater than previous volumetric displays, a dedicated graphics processing engine, and a software subsystem that is easily used with existing applications.

2. SYSTEM DESCRIPTION

Each assembly comprising the 3-D display will now be discussed. Figure 2, above, illustrates a blown-up view of the 3-D display with several components not shown for simplicity. An XGA-resolution (1024 x 768) projection engine illuminates a diffuse projection screen that rotates with projection optics and three relay mirrors at or above 600 rpm. The projected images are radially-disposed "slices" through a 3-D dataset, which when projected in rapid sequence are perceptually fused into a sharp, volume-filling three-dimensional image. At the time of writing, the prototype enclosure, optics, mechanics, rendering algorithms, and graphics processor have been constructed and tested; various electronic subsystems are in the final stages of debugging. The system is designed to provide the following performance specifications:

- At least 216 radial image slices, each of 768 x 768 resolution. System memory sufficient to store imagery comprised of 500 slices that may be used with sufficiently fast projector systems. A minimum of 127 Mvoxels are electronically-addressable; clipping algorithms reduce this to approximately 90 million "API-addressable" voxels.
- 3-bit color (8 colors)
- Viewing angle: 360 degrees horizontal, 180 degrees vertical
- 10"-diameter roughly spherical image volume. Prototype base is 24" wide. Top of dome is 21" from bottom of base.
- No head-tracking or additional eyewear required
- Minimum 20 Hz volume refresh rate
- Application programming interface (API) similar to the OpenGL® API.

A high-level system overview is illustrated in Figure 3, in which data are acquired from the user's computer (the "Host PC") and sent to the volumetric 3-D display. The display contains a custom graphics processing architecture that converts the data into the appropriate series of projection patterns in a process known as *rasterization*.

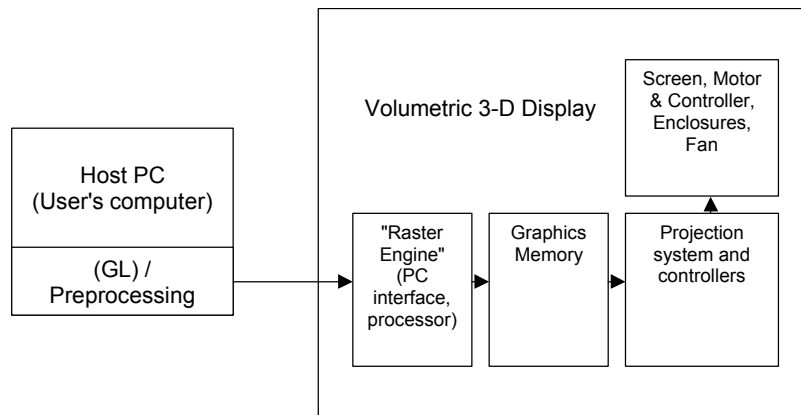


Figure 3. High-level system architecture. A "Host PC" sends 3-D geometry data or volume data to the 3-D display peripheral which performs rasterization, image data buffering, and display.

2.1 Optical / Mechanical Assembly

2.1.1 Design Objective

The hardware is designed to provide high-resolution, high-brightness imagery onto a rotating screen while maintaining a reasonable package size. The packaging objective was to design the 3-D display such that it was approximately the same size and weight as a 21" monitor. The image also had to be at a comfortable viewing height while sitting at a standard workstation. This presented a challenging packaging problem because the projection system's magnification requires long path lengths. Several iterations yielded a tightly-integrated prototype that can be manufactured easily.

2.1.2 System Hardware

An image source was chosen that is capable of providing a high frame rate so that the display will have a minimum of 216 image slices at a 20 Hz volume refresh rate, or just over 4 kHz. Few spatial light modulators exist that are capable of producing XGA-resolution (1024 x 768) imagery at this frequency. Actuality's prototype displays are constructed with Texas Instruments DMD™-based projectors, which utilize MEMS reflective pixels that may be updated at high frame rates. A TI three-DMD™ projection engine is illuminated with an UHP mercury discharge lamp as shown schematically in Figure 4 below. Light from the UHP lamp passes through an integrator which homogenizes the beam and is then shaped through a series of condenser lenses. A cold mirror folds the energy into the prism assembly. The prism consists of a series of elements that separate the red, green, and blue wavelengths to each DMD™. Images formed on each DMD™ are recombined in the prism assembly and are projected to the screen through a series of mirrors and lenses. The image is projected through the center of an open frame DC motor that rotates the three final fold mirrors and the screen. As mentioned above, the projected image rotates in the plane of the screen as the screen assembly rotates. This effect is corrected in software, as discussed below. Access to the projector pixels is clipped to a central region measuring 768 x 768 pixels.

The projection lens system has been designed to minimize axial color shift and field curvature. Several optics are used to optimize the image quality at the screen. The net image also has low astigmatism and fairly low lateral color shift. Distortion at the image is in the form of straight keystone that is corrected in software. The residual pincushion distortion is < 2%.

Image quality at the image plane is sufficient to resolve the 40 lp/mm limiting resolution of the DMD™. At the image plane this is equivalent to 2 lp/mm. Compensation for manufacturing variation within tolerance is accomplished by changing the airspace between the display's relay system and field lenses. A photograph of the actual system is shown in Figure 5.

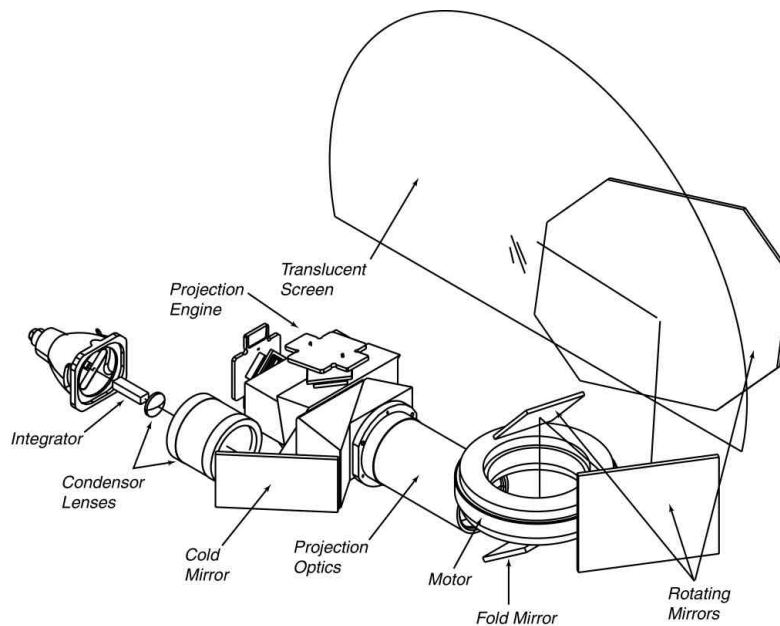


Figure 4. Schematic layout of projection optics assembly – some components not shown for simplicity.

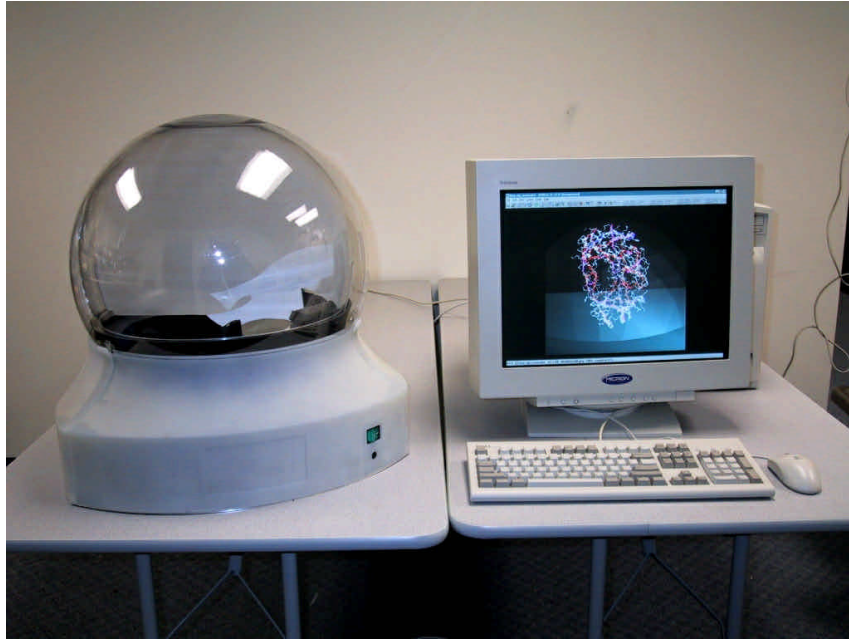


Figure 5. Actual system assembly.

2.2 Rasterization Hardware

2.2.1 System Electronics

Although some pre-processing of geometry information is performed on the host workstation, the rasterization and display electronics are completely contained inside the display. Figure 6 shows a block diagram of the system electronics. In a process described in section 2.3, the Raster Engine accepts volumetric or geometry-based 3-D data from the host PC, converts it into voxel data in the 3-D-display-specific coordinate system, and writes it into graphics memory. The graphics memory module buffers the entire volume of data for continuous display of three colors, and the DMD™ formatter electronics convert the LVTTTL data and control signals into the non-standard voltages needed to operate each DMD™. Screen speed and position are locked to the volume scanning signals by the motor controller.

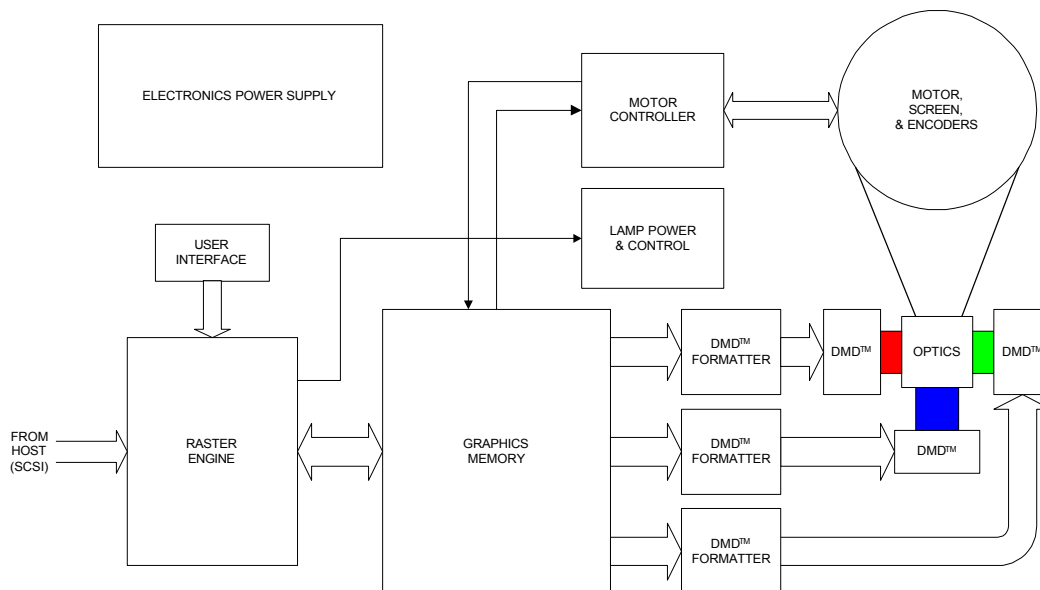


Figure 6. System Electronics

2.2.2 Raster Engine

Computations are performed on a Texas Instruments™ TMS320C6201 32-bit, fixed point digital signal processor (DSP). Executable code, computational tables, and FPGA code are stored in synchronous burst flash ROMs permitting field upgrades over the SCSI interface. Synchronous dynamic RAM (SDRAM) is used for storage of active geometry information and scratch space for the computations that don't fit in the data cache of the DSP. The DSP interfaces to a SCSI controller through a field programmable gate array (FPGA) with a PCI interface. Geometry information can be transferred directly from the SCSI controller to SDRAM by SDRAM and DMA controllers also implemented in the FPGA. This greatly improves the computational efficiency of the rasterization code when operating directly from the program and data caches.

2.2.3 Graphics Memory Module

The 6 Gbits of Double Data Rate (DDR) SDRAM in the graphics memory modules is divided into three colors and two buffers. Scanning of an "active" buffer, random reads and writes to an "inactive" buffer, and refreshes are scheduled and controlled by a graphics memory manager. Data are routed between the CPU, graphics memory, and DMDs™ by three voxel routers. See Figure 7 for a block diagram of the graphics memory module.

Because updating graphics information often involves logic operations between new data and data already being displayed, an automatic read-modify-write capability has been designed into the voxel routers. The Raster Engine can initiate each read-modify-write operation with only one write instruction. AND, OR, XOR, and masking operations have been implemented. Finally, all instructions are buffered in a FIFO in the FPGA to be performed when scheduling allows. This contributes to very low overhead for the Raster Engine software.

The graphics memory manager and each voxel router are implemented in FPGAs. Like the FPGA on the Raster Engine, the code for these FPGAs is stored in the flash ROM on the Raster Engine, permitting field upgrades.

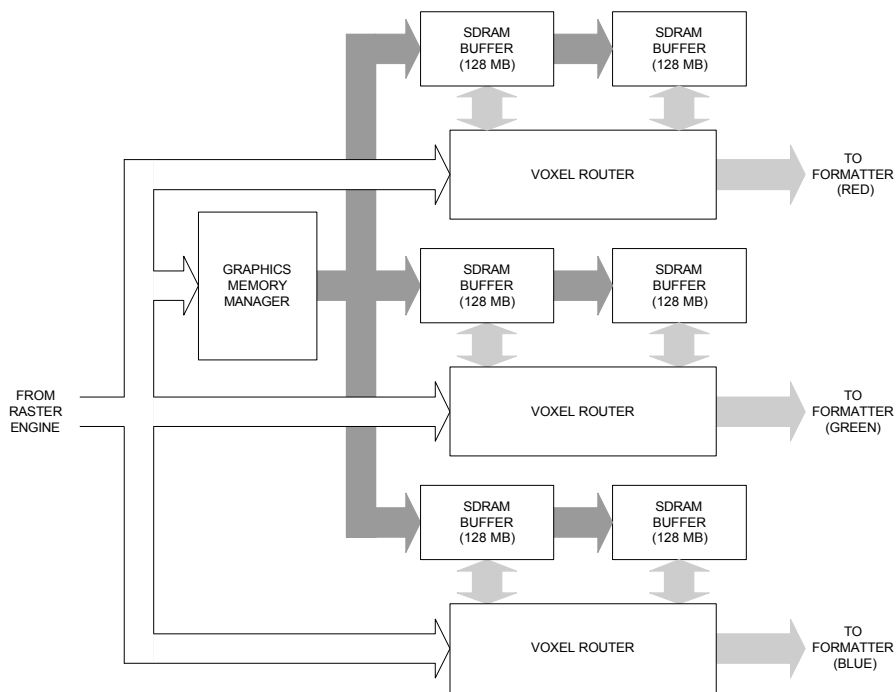


Figure 7. Graphics Memory Module

2.2.4 Motor Controller

The screen position is monitored with four optical encoders whose output is phase-locked to reference clocks generated by the graphics memory manager. A complex programmable logic device (CPLD) computes the speed and phase error between these sets of signals. A low resolution DAC and an analog proportional-integral-derivative (PID) controller provide the necessary control compensation. Pulse width modulation and commutation for the system's brushless DC motor are generated by the same CPLD as well as a soft start for the motor to prevent current surges during power-on.

2.2.5 System Performance

The ability to draw triangles or other data into the display volume will be limited by either the bandwidth of the data paths or the speed with which the voxels can be rasterized (computed by the DSP). The limitation of the SCSI-2 (or Ultra-SCSI) interface is 20-40MB/sec depending on the SCSI port on the host workstation. Since this data usually consists of geometry information, its relationship to voxel bandwidth depends on the size of the polygons being described.

The capacity of the interface between the Raster Engine and the graphics memory module is approximately 133 MB/sec. Voxels may be written one at a time, or as 32-bit words of consecutive voxels. The type of data being displayed dictates which type of write is more efficient.

To maintain a volume refresh rate of 20 Hz, data must be scanned out of the graphics memory at a very high rate. The DMD™ has 1024x768 pixels that must be written for each slice, and there are a minimum of 216 slices. There are also three colors being displayed simultaneously. $3 \text{ colors} \times 216 \text{ slices/color} \times 786,432 \text{ voxels/slice} \times 20 \text{ Hz} = 10.2 \times 10^9 \text{ voxels/sec}$. Since there is one bit per voxel, this corresponds to a bandwidth of 1.3 GB/sec.

2.3 Software

The volumetric display system includes several software components that work together to take data from various applications and render them as 3-D images. The display's native API is based on the OpenGL® API standard. Most features of the OpenGL API are implemented for the 3-D display, and special controls are added as extensions. Several unmodified off-the-shelf applications have been tested and work properly with a software-based 3-D display simulator. Since the display's API is based on the industry-standard OpenGL API, many existing applications can use the display without modification as long as hidden surface removal is not performed before OpenGL in the rendering pipeline. Additionally, programmers can write custom software to take full advantage of the volumetric display without learning a new graphics language.

Actuality's 3-D display graphics API is built on top of Mesa, an open-source OpenGL® API work-alike library. The library closely follows the OpenGL specification, so that 3-D applications that run on the host computer with a 2-D display will run in the same manner with our volumetric display. The display's API replaces the original OpenGL API on the host computer. When the 3-D application calls the OpenGL library to render an image, our Mesa-based library is called instead. All graphics calls are normally passed through our library to the host computer's native graphics library, so that the traditional 2-D projection is rendered unmodified on the host computer's monitor. The library splits off a copy of the function calls it receives to the original OpenGL library, which has been moved to a new location. The original OpenGL library performs the necessary rendering for the 2-D CRT in exactly the same way it did before installation of the volumetric display. That is, the 3-D display augments existing workstation hardware; it allows the user to see graphical output simultaneously on both the traditional CRT and the 3-D display. All the original functionality of the application is maintained, including performance benefits from 2-D graphics hardware when writing to the CRT.

This concept was proven by running off-the-shelf applications for molecular visualization and mechanical CAD using this method for acquiring calls to OpenGL. We ran an unaltered version of a popular MCAD package, sent a copy of the GL calls over the office LAN, and rotated the imagery on another computer which ran a software-based simulation of the 3-D display. Please see Figures 8 and 9. The 3-D display simulator was particularly useful for testing the novel line- and triangle-drawing algorithms required by the display.

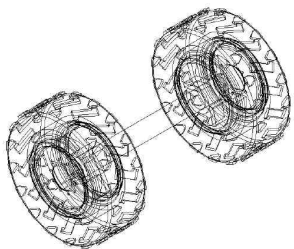


Figure 8. Screenshot of commercial MCAD application.

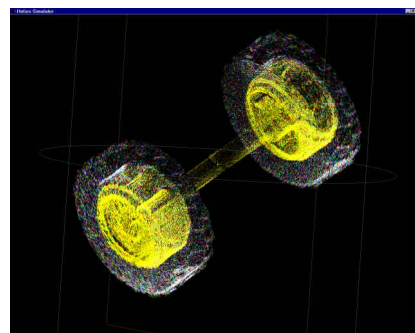


Figure 9. Screenshot of 3-D display simulator.

The Mesa-based OpenGL-like library does the necessary pre-processing on the GL data to prepare it to be rendered by the rasterizing hardware. This involves splitting polygons into triangles and calculating certain parameters required to draw in the display's coordinate system. Polygons outside of the cylindrical volume of the display are clipped. Once the data is prepared, it is passed on to a SCSI device driver in order to be sent to the display.

The host PC's SCSI device driver is similar to other simple device drivers. It buffers the geometric data generated by the graphics library and transmits it to the Raster Engine. The data passes over a SCSI cable to the device. SCSI processor on the volumetric display, along with custom logic, are responsible for making sure the data from the host computer is put into the buffer for processing by the DSP.

As discussed above, the display's Raster Engine converts data from the host PC into a series of "image slices" that are projected in sequence to appear in 3-D on the display. As mentioned in section 1.2, the 3-D display's coordinate space is the composition of a rotating screen with a rotation in the plane of the screen. Furthermore, the projector system introduces keystoneing. Significant research was required to devise a series of patent-pending line- and triangle-drawing algorithms for this coordinate space, since existing algorithms^{18,19,20} are unsuitable for direct application to this display.

The display is also capable of accepting "volumetric data" or data that is already comprised of points in a grid. This kind of data is typical of MRI, 3-D sonogram, and seismic imaging data. In order to show volumetric data on the 3-D display, it is necessary to resample the data so that the correct points in volumetric space may be illuminated. Once the data is resampled, it may be sent to the display for immediate viewing. For volumetric data, due to the fact that little processing is needed, the speed of the SCSI connection is the bottleneck which limits the complexity and update rate of viewable images on the display. In addition to displaying raw volumetric data it is also possible to transform the volumetric data into polygonal surfaces and then render and display them as is common on today's 2-D CRTs. We have demonstrated successful resampling of medical volume datasets in conjunction with the 3-D display simulator described above.

3. FUTURE WORK

At the time of writing, the display architecture described here has been constructed, and electronic assemblies are being debugged. Near-term steps include a full system characterization. In the longer term, we believe that additional work remains to be done in both increasing and decreasing the size of the image volume; some applications call for screens approximately 3 feet in diameter, while others are more appropriate on the scale of several inches. We intend to continue research in the use of dithering algorithms to increase the perceived color gamut.

4. CONCLUSION

The authors have described an 8-color multiplanar volumetric display that is in the final prototype debugging stages. It is expected to provide volume-filling imagery comprised of over 90 million voxels by using an XGA-resolution projector to illuminate a rotating screen with a rapid sequence of 2-D "image slices." A dedicated graphics processing architecture aids in formatting the user's data for the projector system, where it is stored in 6 Gbits of DDR SDRAM. The technology was designed for immediate usability in a commercial environment, so additional research was performed in ensuring that the 3-D display system would function with a well-known graphics library.

ACKNOWLEDGMENTS

The authors wish to thank the conference Chairs for the opportunity to present this work. The authors also wish to thank Shawn Samuel, William V. Baxter III, and Michael Halle for past and ongoing algorithmic guidance, David Oliver at Cusp Development for industrial design services, and Michael Thomas and Jane Bareau of Optical Product Development for optical design. Digital Micromirror Device, DMD, Texas Instruments, and TI are either registered trademarks or trademarks of Texas Instruments Incorporated. SGI and OpenGL are registered trademarks of Silicon Graphics, Inc. Significant portions of the work described here are patent-pending.

REFERENCES

1. G. Favalora, D. M. Hall, M. Giovinco, J. Napoli, and R. K. Dorval, "A Multi-Megavoxel 3-D Display System for Distributed Collaboration," *IEEE Globecom 2000 workshop, "Application of Virtual Reality Technologies for Future Telecommunication System,"* (San Francisco, CA), 2000.
2. G. Favalora, "Goggle-Free Volumetric 3-D Display of Protein Structure," *IBC's Annual Symposium on Functional Genomics*, (Litchfield Park, AZ), 2000.

3. B. G. Blundell and A. J. Schwarz, *Volumetric Three-Dimensional Display Systems*, p. 4, John Wiley & Sons, New York, 2000.
4. R. D. Ketchpel, "Three-Dimensional Cathode Ray Tube," U.S. Patent No. 3,140,415, 1964.
5. P. Soltan, J. Trias, W. Dahlke, M. Lasher, and M. McDonald, "Laser-Based 3-D Volumetric Display System: The Improved Second Generation," *Naval Engineers Journal*, pp. 233-243, May, 1995.
6. B. G. Blundell, A. J. Schwarz, D. K. Horrell, "Cathode ray sphere: a prototype system to display volumetric images," *Optical Engineering*, **33**(1), pp. 180-186, Jan. 1994.
7. M. Halle, "Autostereoscopic displays and computer graphics," *Computer Graphics*, ACM SIGGRAPH, **31**(2), pp. 58-62, May 1997.
8. K. Langhans, D. Bezecny, D. Homann, D. Bahr, C. Vogt, C. Blohm, and K.-H. Scharschmidt, "New portable FELIX 3D display," in *Projection Displays IV*, Ming. H. Wu, Editor, Proceedings of SPIE Vol. 3296, pp. 204-216, 1998.
9. J. D. Lewis, C. M. Verber, and R. B. McGhee, "A True Three-Dimensional Display," *IEEE Trans. Electron Devices*, **ED-18**, NO-9, pp. 724-732, 1971.
10. E. Downing, L. Hesselink, J. Ralston, and R. Macfarlane, "A Three-Color, Solid-State, Three-Dimensional Display," *Science*, **273**, pp. 1185-1189, 1996.
11. I. I. Kim, E. Korevaar, and H. Hakakha, "Three-dimensional volumetric display in rubidium vapor," in *Projection Displays II*, Proceedings of SPIE Vol. 2650, pp. 274-284, 1996.
12. A. C. Traub, "Stereoscopic display using rapid varifocal mirror oscillations," *Appl. Opt.*, **6** (6), pp. 1085-1087, 1967.
13. M. Hirsch, "Three dimensional display apparatus," U.S. Patent No. 2,967,905, 1961.
14. Author unknown, "New Display Gives Realistic 3-D Effect," *Aviation Week*, pp. 66-67, October 31, 1960.
15. R. G. Batchko, "Rotating flat screen fully addressable volume display system," U.S. Patent No. 5,148,310, 1992.
16. C.-C. Tsao, Z.-H. Zhou, and A. K. Hu, "Method and apparatus for displaying three-dimensional volumetric images," U.S. Patent No. 5,754,147, 1998.
17. C.-C. Tsao and J. Chen, "Volumetric display by Moving Screen Projection with a fast-switching spatial light modulator," in *Projection Displays IV*, Ming H. Wu, Editor, Proceedings of SPIE Vol. 3296, pp. 191-197, 1998.
18. J. E. Bresenham, "Algorithm for Computer Control of a Digital Plotter," *IBM System Journal*, **4**, pp. 25-30, 1965.
19. A. E. Kaufmann, "Method of converting continuous three-dimensional geometrical representations of polygonal objects into discrete three-dimensional voxel-based representations thereof within a three-dimensional voxel-based system," U.S. Patent No. 4,987,554, 1991.
20. T. Inoue, T. Asaka, and Y. Yamaguchi, "Method and device for generation of quadratic curve signal," U.S. Patent No. 4,484,298, 1984.